Creating Knowledge Maps by Exploiting Dependent Relationships

John L.Gordon
PUBLISHED: ES99 Conference – Applications stream. pages 63 - 78,1999 ISBN 1-85233-230-1 **ALSO PUBLISHED** in: Knowledge Based Systems, Vol13 (April 2000) pages 71 - 79, Elsevier Science
NWAIAG, Blackburn College, Feilden Street, Blackburn, Lancashire BB2 1LH.

Abstract

Knowledge is an interesting concept that has attracted the attention of philosophers for thousands of years. In more recent times, researchers have investigated knowledge in a more applied way with the chief aim of bringing knowledge to life in machines. Artificial Intelligence has provided some degree of rigour to the study of knowledge and Expert Systems are able to use knowledge to solve problems and answer questions.

Current business, social, political and technological pressures have forced organisations to take greater control of the knowledge asset. Software suppliers and others, offering valuable solutions in this area have unfortunately clouded the issue of knowledge. Information and data control are seen as implicit knowledge management tools and many have abandoned the search for explicit knowledge management methods.

Knowledge representation schemes help to identify knowledge. They allow for human understanding and machine application and they can support the automated use of knowledge in problem solving. Some of these representation methods also employ spatial techniques that add an extra dimension to human understanding. Knowledge mapping defined in this work uses learning dependency to organise the map and draws on the ideas of what knowledge is and on spatial representation structures. Knowledge maps can support metrics that provide information about the knowledge asset. Knowledge maps create a visible knowledge framework that supports the explicit management of knowledge by organisation managers and directors. Knowledge maps also offer other advantages to the organisation, the individual and to educational institutions.

Contents

Acknowledgements

Thanks to the following people for particular support.
To staff from British Aerospace, Samlesbury, for valuable contributions.
To Brian Whitehead from Blackburn College for help in producing the final draft.
This work is being partly supported by an **ESF ADAPT project.**

1. Introduction

There is little doubt that knowledge is a complex concept that has occupied the thoughts of Philosophers and others for hundreds (thousands) of years. It is not surprising then that the current thoughts on Knowledge Management and efforts to establish such ideas in business and industry, can be difficult or can be inappropriate.

Once people start seriously discussing the knowledge asset of a company, the issue of what knowledge is and if information or data is knowledge soon emerges. Looking at articles from some of the new business magazines concerning knowledge management can lead the reader to imagine that knowledge management is a term that has little to do with knowledge. One article (of many) even stated that knowledge is impossible to manage but managing information (with the author's company's software) has the side effect of managing knowledge. Such statements confuse the whole area of knowledge management and generally have a commercial motive. There is a growing effort to develop ontologys that will help to clarify the area of applied knowledge (ontology 1999).

1.1story of knowledge

Table 1 provides a very brief and somewhat incomplete look at the emergence of ideas about human knowledge. The table shows that knowledge has interested human kind for centuries. Early thoughts considered how our knowledge is derived from our senses. It was later realised that what we sense is not necessarily what actually happens. Two opposing views emerged, the view that knowledge is mainly derived from the world we live in, through experiences, and the view that true knowledge can only be derived from abstract thought. (Epistemology 1999)

These opposing views gradually grew together because philosophers appreciated the contributions made from experience and from abstract thought. Later, other components of knowledge were identified, including tacit knowledge.

Philosopher	Period	Classification	Summary
Georgias	485-380bc	Sophist	Nothing Exists. If anything does exist it cannot be known. If anything exists and can be known it cannot be communicated.
Protagoras	480-411bc	Sophist	Nothing is absolutely good or bad, true or false, so each individual is therefore his own final authority.
Socrates	470-399bc	Rationalist	Every person has innate knowledge of ultimate truth and need only be spurred into conscious reflection to become aware of it. The philosophers task is to provoke thought not to teach. Knowledge originates in sensory perception.
Plato	428-347bc	Rationalist	Reality lies in abstract thought. Abstract knowledge is superior to imperfect concrete observation.
Aristotle	384-322bc	Empiricist	Knowledge is acquired through empirical evidence obtained through experience and observation. Induction of principles from observation. The science of logic represented by the syllogism.
Aquinas	1225-1274		Perception is the starting point for knowledge and logic and is the intellectual procedure for arriving at reliable knowledge. (also believed in faith)
Bacon	1561-1620	Empiricist	First to formulate rules of inductive inference. Called for new scientific method based on inductive generalisation.
Descartes	1596-1650	Rationalist	Based on mathematical proof. Application of deductive and analytical methods.
Locke	1632-1704	Empiricist	Argued that knowledge is derived from experience either of the external world through sensation or the mind through reflection. One cannot have absolute certain knowledge of the physical world.
Berkeley	1685-1753	Empiricist	The only things that one can observe are ones own sensations and these are in the mind. Knowledge comes from ideas but there is no distinction between ideas and objects.
Hume	1711-1776	Empiricist	Knowledge is of two kinds. 1) Knowledge of mathematics and logic which is certain but provides no information about the outside world. 2) Knowledge derived from the senses which is largely a knowledge of cause and effect, which

			means that one cannot hope to predict scientific development or for scientific knowledge to remain true.
Kant	1724-1804		One can have certain knowledge but such knowledge is more informative about the structure of thought than about the outside world. Three types of knowledge: 1) Analytical base truths (uninformative). 2) Synthetic - learned from experience - prone to error. 3) Synthetic base truths - pure intuition (mathematics and philosophy).
Hegel	1770-1831	Rationalist	Revival. Thought and History.
Husserl	1859-1938	Phenomenology	To distinguish the way things appear from the way one thinks that they really are. Understanding the conceptual foundations of knowledge.
Wittgenstein	1889-1951	Logic empiricism and positivism.	Use of language. Tacit knowledge

Table 1

Computer based Expert Systems and Artificial Intelligence have also made an important contribution to our understanding of knowledge. For this activity to succeed, researchers had to be very clear about what they meant by knowledge and had to develop rigorous representations for knowledge so that the knowledge could be brought to life in a computer program (Shortliffe 1976). The contribution of Artificial Intelligence to the understanding of knowledge has been significant. Section 2 looks a little more closely at some of these more rigorous knowledge representation methods.

1.2 Human and Machine Knowledge

Knowledge is a subject for debate and analysis in a broad range of disciplines from the philosophical basis of Epistemology to the application based Knowledge Engineering used in Artificial Intelligence. Within all of this work, precise definitions are still elusive and the meaning of knowledge is largely relative. There are however some clear distinctions between machine knowledge and human knowledge. Whether these distinctions will remain true is more doubtful.

Although dictionaries may have difficulty separating the words ‘know’ and ‘understand’, if a computer based Expert System can be said to know things or have knowledge, it is not possible to extend this to say that the computer understands. For a computer to function in an expert domain, it is given the necessary knowledge and is able to use this knowledge to solve problems or give advice. When a human expert possesses the same knowledge, we may withdraw the status of expert if that human is found not to understand the knowledge, but simply to believe it to be true. This is a significant difference between knowledge which we say may be possessed by machines and knowledge which is possessed by humans. It may be an oversimplification to say that ‘Intelligent machines’ ‘know’ and ‘Humans’ understand, but this is a useful generalisation.

Why is there understanding when humans acquire knowledge?

Most understanding comes from a deep and rich knowledge, from an ability to work out why. Knowledge is learned incrementally, some things need to be learned before it is possible to learn other things. That is not to say that humans could not function just like computer based expert systems and use knowledge without understanding it. Normally however, human experts do understand and have acquired knowledge through an incremental process which leads to the acquisition of the expert knowledge in question.

1.3 The increasing role of knowledge

We are now in, or at least on the edge of, an information age. Information is being used as a commodity and many companies exist in the information sector. Data is now more widely available due to great improvements in capture and storage of data. More data is now being derived automatically through new sensing techniques. This means that any bottleneck caused by relatively slow data entry from humans has been overcome (in many situations). In fact, many authors have written about the explosion of data and information that society is now experiencing. (Frawley 1991)

Artificial Intelligence has responded to the perceived challenge by developing new tools to produce knowledge from data (Hutchinson A. 1994, Thornton C.J. 1992). Such tools exist along with the assumption that the transformation from data (or information) to knowledge can be clearly stated and well understood. Indeed, for these systems and within their domain, the transformation (or derivation) is well understood. This also implies that the distinction between knowledge and information or data is clear. For Artificial Intelligence, this clear distinction does exist. However, the distinction is not always relevant to all areas of knowledge and information, particularly the area of tacit knowledge. Within the domain of knowledge management, the clarity has almost disappeared.

1.4 Managing the invisible

Knowledge is a complex concept and is itself, invisible. These two factors lead to difficulties for those attempting to manage knowledge. One of the more serious problems is that informed opinion between managers based on common experience and representation is difficult. People have their own views concerning knowledge. These views may be difficult to communicate meaningfully to others. Discussion on this basis may involve one person trying to describe an abstract concept to others before using this concept to develop an argument. This is unlikely to result in a common view of issues about the knowledge asset and may even result in misunderstanding about specific actions to be taken.

There is a need for a way of helping people to visualise knowledge and maintain and develop a common visualisation and representation. Successful representations have already been developed but these are more suited to computer based processing. Even so, these representations do bring clarity to the situation and can be valuable in the human as well as the machine domain.

2. Knowledge Representation

There are several accepted methods of knowledge representation that have been devised for AI type applications. Some of these are also suitable for use and interpretation by humans and can form a bridge between human knowledge and machine knowledge. This is important if organisational knowledge is to be archived in such a way that it can be effectively used in automated systems and also understood and updated by humans. Some of these representational methods will be discussed in order to allow the reader to consider their merit as knowledge management tools.

2.1 Rules

Rules are reasonably easy to understand by humans and are also a powerful machine based knowledge representation scheme. Rule based systems that could apply human knowledge and function at the level of a human expert were famously pioneered by E.H. Shortliffe in the system 'Mycin' (Shortliffe 1976). Rules require knowledge to be identified as attribute value pairs. They take the general form:

if attribute A1 has value V1
and attribute A2 has value V2
then attribute A3 has value V3

Attributes can represent internal data items, they can represent input or output systems or they can initiate a response from the user. Once knowledge is represented as a rule set, it is relatively easy to construct an engine that can make use of the rules in an automated reasoning system.

Attributes can represent internal data items, they can represent input or output systems or they can initiate a response from the user. Once knowledge is represented as a rule set, it is relatively easy to construct an engine that can make use of the rules in an automated reasoning system.

In addition, the rules themselves can be archived and updated as necessary. This would be a knowledge archive rather than an information archive since the rules can be directly used in automated reasoning.

Exception systems are similar to rules in that they can also be archived, understood by humans and used directly in automated reasoning systems. Exceptions may take the form:

attribute A1 has value V1
unless attribute A2 has value V2
and attribute A3 has value V3

2.2 Frames

Frames are also a powerful knowledge representation system that are accessible to both humans and machine. A frame is a collection of information and associated actions that represents a simple concept. It would be possible to represent a person (in a simple way) by the use of a frame.

Frame:	Elery Stone	
	Specialisation of:	Frame Person
	Date of Birth:	30:04:62
	Sex:	Male
	Nationality:	British
	Home Town:	St. Helens
	Occupation:	Tailor
	Health:	(Consult Medical system)

In the simple frame shown above, most of the slots have values but one slot requires an independent system to be called to find a value. Frames are a mixture of information, calls to information derivation functions and output assignment. Frames can be used to represent complex pieces of knowledge and can also be archived and edited as required.

2.3 Semantic Networks

Semantic networks are a powerful knowledge representation system. They are easy to understand by humans and can be used in automated processing systems. This means that they can also become a vehicle to archive company knowledge. A typical semantic network that represents knowledge concerning an electric space heater could be:

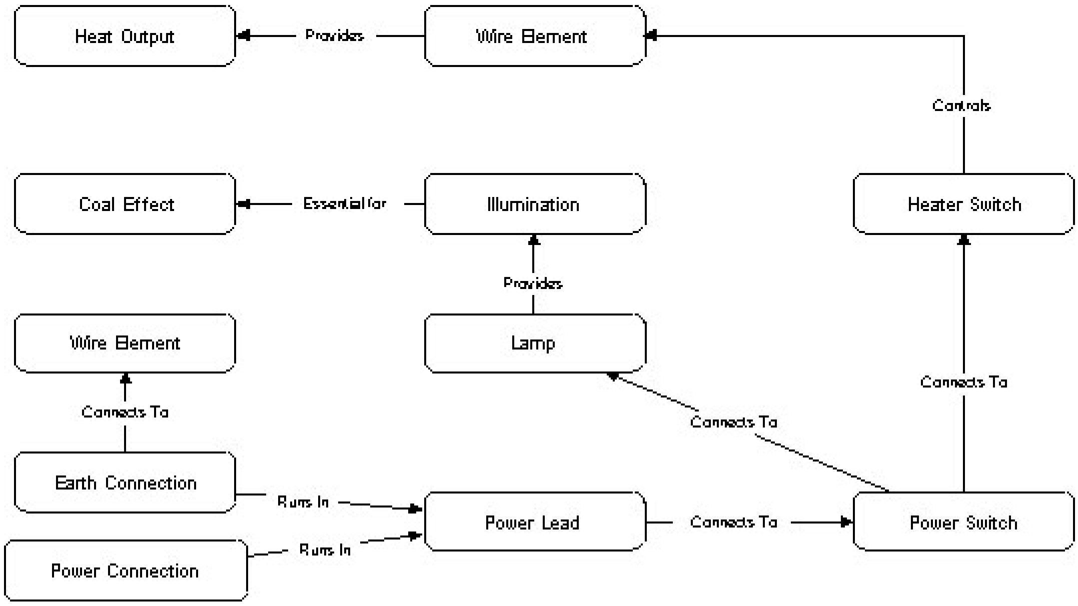

Figure 1: Semantic Network to describe an Electric Heater

In this simple network, nodes are specific items and links show relationships between items.
It would be possible for an automated system to answer questions about items contained within the network by following links (provided that it could understand the questions).
How does power get to the heating element?
What is the purpose of the lamp?
It would also be possible to a computer to construct a textual statement about the knowledge contained in the network.

2.4 Concept Diagrams

Concept diagrams are closely related to semantic networks. Concept Diagrams are also composed of nodes and arcs and the nodes and arcs have similar functions. Concept diagrams can be used to describe fairly complex concepts and are suitable for both machine and human interpretation. They are seen as a knowledge representational method that employs graphical structures (Sowa 1984). There is a body of work relating to concept diagrams and their use as a graphical logic. (Sowa 1993). This offers interesting opportunities for work on knowledge mapping by creating the framework that could allow knowledge maps to be transformed into other machine understandable representations such as the Knowledge Interchange Format (KIF) (Genesereth 1992).

3. Structural Representation

The diagrammatic knowledge representation methods described in sections 2.3 and 2.4 are not only suitable machine representations but provide a more appropriate representation for human understanding because they include spatial as well as textual information. This sort of diagram allows groups of people to share a common understanding of a complex topic. This is the sort of representation that is appropriate for knowledge structure representation.

3.1 Knowledge Structure

Within the context of this work, the actual knowledge is not directly part of the structure of knowledge but is indexed from it. In order to create a structure for knowledge , it is necessary to identify specific pieces or islands of knowledge and give them a unique name or identifier. These identifiers can then form part of a structural diagram for knowledge and can also be used to index to the actual knowledge implied by the identifier. The amount of knowledge that an identifier represents, or granularity, is an important consideration but should match the context within which the diagram will be used. If the knowledge concerned with boiling an egg is considered, then the associated identifiers may be.

Boiling an egg
Boiling Water
Obtaining an egg
Chemical changes

Each of these identifiers can represent a piece of knowledge but are not that knowledge. Similarly, the knowledge concerning the calculation of the gravitational attraction between two masses may include:

Gravitational Attraction
Mathematics
Mass
Force
Distance

Once knowledge identifiers can be derived, it is then necessary to consider a relationship between these identifiers that is both valid and useful within the context of knowledge mapping. At the present time, the acquisition of machine knowledge employs little dependent structure but human knowledge, particularly expert human knowledge, is acquired in a more rigorous way. Human knowledge is learned and the learning process is dependent on prior knowledge. That is, the learning process is hierarchical because the understanding of new knowledge often relies on the prior understanding of some existing knowledge.

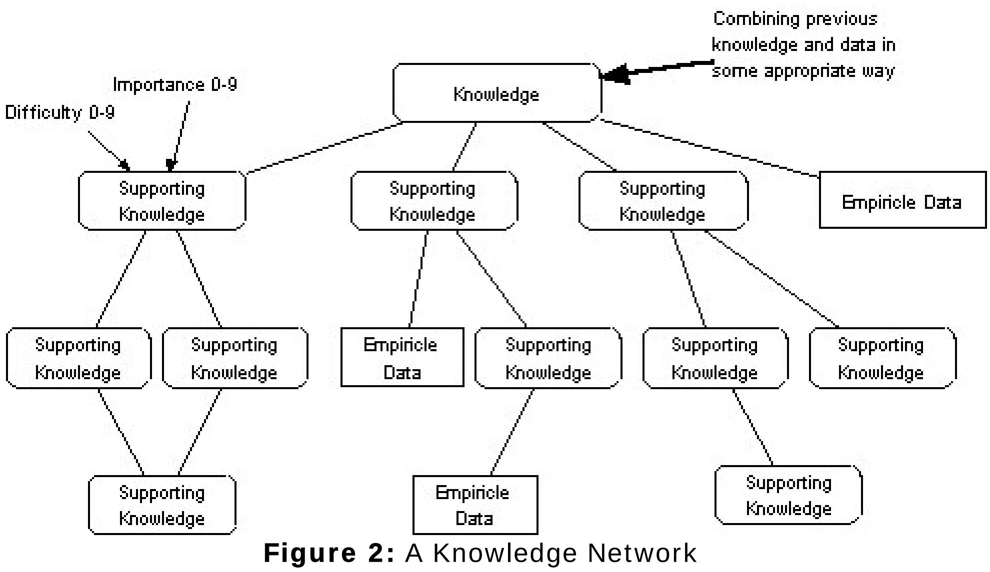

Figure 2: A Knowledge Network

Figure 2 shows what a knowledge network may look like. Understanding implies that the human expert knows why the supporting knowledge or empirical data is actually supportive of the higher knowledge item. Also contained within each link would be a measure of importance which shows how important each supportive piece of knowledge or empirical data is to the higher knowledge item.

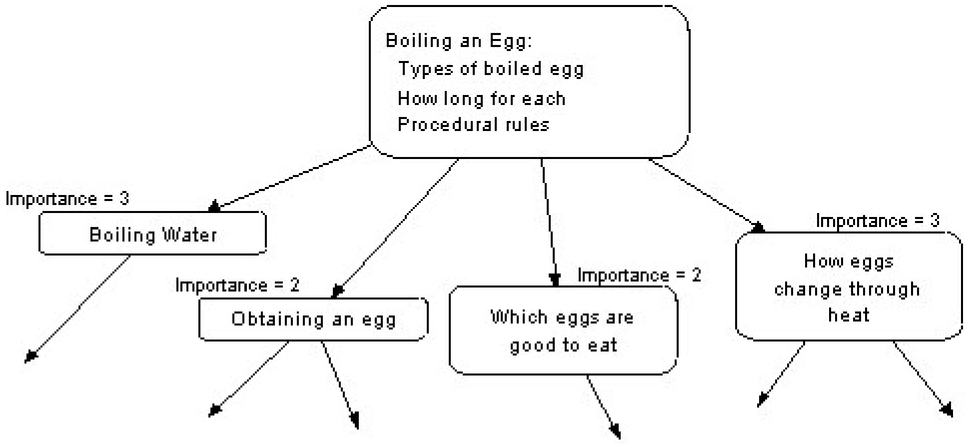

Figure 3: Boiling an Egg

An incomplete example is shown above; that of boiling an egg. The knowledge concerning boiling an egg involves knowing about hard and soft boiled eggs and how long it takes to produce each. It also involves knowing the procedure and procedural rules for completing the operation. This in turn implies that the expert egg boiler already knows how to boil water, which is clearly an important part in egg boiling because it is water that is used to boil eggs in. Knowing how the inside of an egg changes with boiling is also very important if the expert is to understand what makes a cooked, rather than a raw egg and what causes the change to take place. Almost as important is obtaining the egg in the first place and which eggs are normally eaten. Clearly egg size would also affect the procedural rules.

This example illustrates the point that an expert egg boiler can do more and knows more than simply how to boil an egg. A robot egg boiler however, may be able to boil adequate eggs but may abstain from engaging in a conversation about egg boiling.

Another example from a different class of knowledge could be knowledge concerning gravitational attraction.

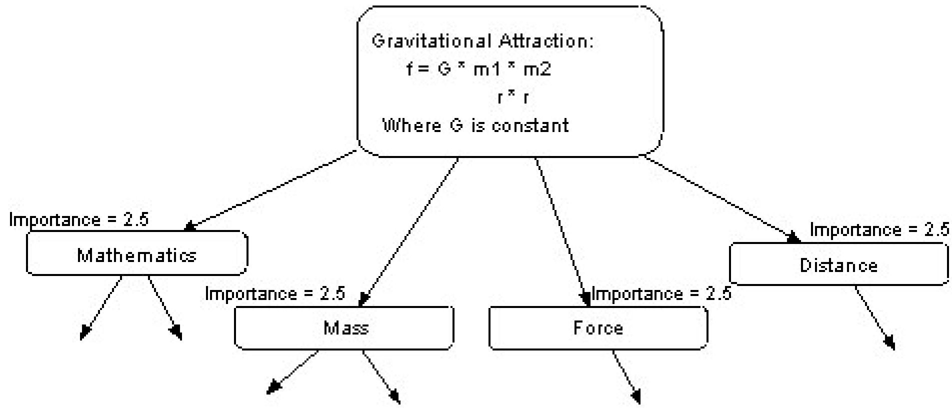

Figure 4: Gravitational Attraction

Again, no attempt has been made here to expand the network to include the knowledge necessary to know 'mass' for instance.
Looking at expert knowledge in the case of 'gravitational attraction' leads to a similar conclusion to that derived for 'egg boiling'. That is, a computer which exhibits knowledge relating to gravitational attraction could be expected to be able to calculate force, to know what if any parameters are required in order to calculate force and to provide simple explanations of what gravitational attraction is. A human could also do this given the formulae and instruction on how to use it and some answers to specific questions. However a human who can be said to understand gravitational attraction will have a clear understanding of concepts such as force and mass. That human would also understand other concepts that are the prerequisites of mass and force. Therefore, the human expert

is much more likely to be able to answer novel questions about gravitational attraction and also see gravitational attraction in new and interesting situations.

3.3 Knowledge Maps

The framework derived in section 3.2 forms the basic framework for the knowledge map as identified in this work. Unfortunately, the term 'knowledge map' is far too general for this to be the only derivation. It is certain that others will have been and will be derived. 'Knowledge map', as used in this work relies on the underlying principle of learning dependency or prerequisite knowledge. The idea has been explored in other knowledge management tools and various methods of visualising the map have been briefly discussed (Gordon 1997).

In terms of application, this derivation of the knowledge map is highly relevant to a range of human endeavours. This will be explored further in section 4. It is the relevance to current work surrounding knowledge management and knowledge based systems that provides some justification for the use of the term knowledge map.

There is some question as to an appropriate granularity for a knowledge map. In principle, a knowledge map would contain very elementary pieces of knowledge and therefore, even a simple knowledge map would be very large (because of high granularity, many nodes). Any practical use of knowledge maps would contain flexible granularity. Some pieces of knowledge may themselves consist of knowledge maps that show a greater granularity. It would therefore be possible to explode a knowledge map containing say 10 items, into a higher granularity of say 100 items. In the latter case, the connectivity would also become more complex. Therefore, the granularity of a knowledge map should be as low (few nodes) as possible whilst still providing all of the functionality for which it was intended.

4. Applications & Case Studies

The discussion in this section is aimed at both potential and actual applications of the concept of knowledge maps based on learning dependency. The first application to be discussed in general terms involves some work carried out with British Aerospace at Samlesbury, concerning the mapping of knowledge in a particularly sensitive area of the business, using the framework derived in this paper.

4.1 Knowledge Mapping at British Aerospace

This particular project was initiated by British Aerospace for two reasons. The first was that they had identified a specific problem concerning expert level knowledge that was being shared with other organisations. The problem was that although the company had developed and pioneered this knowledge, they did not know exactly what was being shared and therefore had no way of knowing the value of what they were essentially giving away. The second reason is that British Aerospace was aware of and involved in the research being carried out at the NWAIAG concerning knowledge mapping.

The actual project was of a pilot, investigative nature. Two people spent one week at the factory, following a prescribed methodology, talking to and interviewing staff as well as touring the site. The site was already familiar to the project staff. The objective was to create a large (A0), printed map of the knowledge contained within the specific area in question (as far as time would allow) and to add some qualitative measures to the elements contained on the map. The qualitative measures were estimates of the importance and difficulty (to learn) of each knowledge element and if the knowledge was more procedural or declarative in nature. In this instance, the metrics were attached to the nodes and not to the arcs. The investigation also involved a report from a short tour of the facility, the sort of tour a visitor may be taken on.

The investigation concentrated on areas of knowledge. Inevitably, the staff being interviewed were happier talking about processes rather than knowledge. However it was possible to get them to talk about the knowledge required to successfully carry out certain process activities. It was also possible to get the staff to talk about prerequisite knowledge (what do you need to know before you can learn and understand that?). Knowledge nodes were added to a computer screen as the interview proceeded, in full view of both interviewer and interviewee. Prerequisite knowledge was shown as directional arcs between nodes. Interviews involved both verification and elicitation to ensure the construction of a map that was generally acceptable.

The resulting map showed the items of knowledge that were required in the particular section being investigated. The connectivity showed hierarchical learning dependency and importance and difficulty to learn were colour coded onto each node. This gave a map that could form the basis of management discussion concerning the knowledge used in and required by the section. It showed clusters of knowledge, it identified high risk knowledge and it allowed managers to discuss what part of the knowledge was essential to the company and what part was not. It also allowed training schemes to be discussed by virtue of the knowledge hierarchy. The company estimated the value of the knowledge contained within the work to be at least a six figure sum.

4.2 Knowledge Mapping in Business

The project identified in section 4.1 is not an isolated example. In general, this sort of knowledge mapping is aimed at creating a visible framework for knowledge that will facilitate its manageability. The method is not intended to replace other efforts a company may use to manage its knowledge or information resources. It is however, aimed at supporting managers in their efforts to explicitly manage knowledge rather than create automated solutions that will manage knowledge implicitly.

If knowledge maps, of the type described, were readily available to companies, managers would be able to retain a common view of the knowledge asset and begin to plan schemes to target critical knowledge areas. These may be areas that are essential for the companies survival or they may be areas that have been shown to be high risk (of loss). It would be up to the managers of the companies to develop strategies that matched their budget and that fit in with other medium and long term company plans. It would also allow managers to consider the knowledge asset when planning other strategic business changes.

Other areas that have been investigated as part of this work include prototype design and manufacture and one section in a now private utility. In each case, similar methods were used and similar benefits identified. The work at British Aerospace is the most advanced of these projects.

4.3 Knowledge Mapping for Curriculum Development

Many educational establishments already plan and identify course structure using some sort of progress map. Such maps can show which courses lead to other courses and which courses together represent a specific area of expertise. The maps, although helpful, are very crude (they lack detail) and because of this have major shortcomings. For instance, if a student can pass a course by achieving 45% in an end of term examination and that examination tests about 50% to 60% of the knowledge delivered in the course, it is possible for a student to pass on to the next stage with considerable gaps in his or her knowledge. This may not be important from the view of gaining a qualification but it may be a very important factor in preventing that student from acquiring new knowledge at the next level. There are assumptions made at each level of study regarding prior knowledge. If this prior knowledge is not there, then it may not be the students or the teachers fault, but the student can easily fail at this stage. It may be the system of study that is at fault in not making sure that new knowledge is presented only when prerequisite knowledge has been acquired.

By investigating the knowledge needed in a particular area of study (in a finer way) and then mapping out this knowledge using learning dependency, prior knowledge assumptions will be clear to both student and teacher. Students will be able to see clearly why they must know a piece of knowledge and not simply pass an examination in it. They will also see that knowledge is accumulative and is not simply to be forgotten after testing.

Creating such maps would be a major undertaking but the potential benefits for greater modularisation of learning and support for individual distance learning would be significant.

4.4 Knowledge Mapping for Personal Development

Section 4.3 discussed the potential for this type of knowledge mapping in curriculum design and also suggested benefits for the individual learner. The individual learner could benefit from the availability of a clearer picture of a study programme. The map would show the learner that it is better to master each prerequisite knowledge node before attempting the next level. The learner would also be able to identify and plot a course to a learning goal.

Companies could employ the mapping method with individual staff members as part of a staff development programme. This would be particularly valuable if the personal map had the same spatial structure as the section knowledge map. Managers would be able to identify key individuals and would also be able to plan appropriate and efficient training programmes. Inefficiency is often seen in company training schemes because there is not enough care taken to ensure that all trainees (students) have the necessary prerequisite knowledge to benefit fully from a course.

5 Conclusion

This paper has tried to show that although knowledge can be a complex concept and has a rich and extensive background of philosophical analysis, it is both important to organisations and manageable by them. The blur between knowledge and information need not exist. The study of Artificial Intelligence, Knowledge Based Systems and Knowledge Engineering has provided a great deal of objectivity in the realms of knowledge elicitation, analysis and representation and has also provided definition. This background can support knowledge management in organisations. There is no attempt in this paper to suggest that information management systems have little value or that information mobilisation through data warehousing and access systems is not beneficial. They have been shown to be highly beneficial to organisations. However, these things alone do not wholly manage knowledge. It is suggested that they manage knowledge implicitly and this may be true. However, this does not mean that there are not ways to manage knowledge explicitly.

One of the ways that can support explicit knowledge management is to make the knowledge visible in some real way. This paper has discussed the idea of knowledge mapping using identifiers for distinct pieces of knowledge and using learning dependency as the connective structure for the map. Learning dependency can be shown to be a very useful way of organising human expert level knowledge. The uses apply to organisations and to individuals. It is also possible to elicit information about each piece of knowledge that relates to its importance, difficulty to learn and its learning type. The dependent structure provides organisation for the knowledge map and helps to identify clusters of knowledge. The nodes along with their metrics help map users to identify and target areas that require management attention.

Several pilot trials of the knowledge mapping method, using an implementation formula prepared in advance, have produced output of real and in some cases, quantifiable benefit to organisations. The method supports the organisation in several ways:

* It makes knowledge visible to all managers.
* It helps managers identify areas of knowledge requiring attention.
* It allows knowledge to feature in strategic planning and change.
* It can improve the efficiency of staff development.
* It allows managers to make decisions about the knowledge asset.

The method also has benefits for the individual and for organisations specialising in education:

* It allows an individual to see and understand a development programme.
* It helps to concentrate effort on understanding rather than on passing examinations.
* It can help individuals plan their own learning when working alone.
* It can support the comparison of organisational needs and individual attainment.
* It can help educational institutions to plan more efficient modularisation of education to support greater access.

These claims may seem a little ambitious. They also rely on some investment in the knowledge mapping process. However Artificial Intelligence has provided a rigorous foundation for the work and it is likely that elicitation and representation tools could easily be modified and applied to this work. During the initial studies, we (the NWAIAG) have developed an elicitation tool called SKAT (structural knowledge auditing tool) that supports the interview process. We intend to add greater functionality to this program so that it can perform some analysis and support representational structures. Some experiments have been undertaken where pieces of knowledge were elicited and embedded and indexed from the mapping (auditing) tool. These were presented at a business seminar in March 1999 and received favourable comment. Representational methods included video, animation, text, diagrams, pictures, and rules that were made active by a simple backward chaining expert system shell written in a multimedia tool development language.

References

1. Ontology 1999. http://wings.buffalo.edu/academic/department/philosophy/ontology/.

2. Epistemology 1999. Microsoft Encarta 99

3. Shortliffe E.H. 1976. Computer Based Medical Consultations: MYCIN. Elsevier, New York.

4. Frawley W.J., Piatetsky-Shapiro G, Matheus C.J. 1991 Knowledge Discovery in Databases. MIT Press Pp1,27.

5. Hutchinson A. 1994, Algorithmic Learning. Clarendon Press

6. Thornton C.J. 1992 Techniques in Computational Learning. An Introduction. Chapman & Hall

7. Sowa J.F. 1984. Conceptual Structures: Information Processing in Mind and Machine. Addison Wesley

8. Sowa J.F. 1993. Relating Diagrams to Logic. in Concept Graphs for Knowledge Representation Eds Mineau G.W., Moulin B. Sowa J.F. Springer-Verlag pp1-35

9. Genesereth M.R. Fikes R.E. 1992. Knowledge Interchange Format, Version 3 Reference Manual. Report Logic-92-1, Computer Science Department, Stanford University.

10. Gordon J.L. Edge M. Focused Knowledge Management, 1997, in Applications and Innovations in Expert Systems. Ed. Macintosh A. Milne R. SGES Publications pp207,219